



Association Professionnelle
des Valorisateurs
d'Instruments Financiers

Valuation of Financial Derivatives Practical Guidance

SCOPE

This document intends to give practical guidance for the valuation of financial derivatives which require the use of a model, together with its algorithm implementation, and a set of parameters to produce a theoretical value.

Authors:

Mourad BENALI

Eric BENHAMOU

Francis CORNUT

Frederic DES CLOSIERES

Claude MARTINI

Version:1.0

FINAL DRAFT

Last updated:June 28, 2015

Contents

1	Introduction	2
2	Definitions	3
2.1	Financial instruments	3
2.1.1	Derivative instruments	3
2.1.2	Basic instruments	5
2.1.3	Underlyings	6
2.2	The concept of value	6
2.3	Valuation methodology	8
2.3.1	Do we understand the product for which we must produce a valuation?	8
2.3.2	Risk free rate?	8
2.3.3	Should models define inputs or inputs define models?	9
2.3.4	Fair value or valuation outputs?	9
3	Market data and calibration	10
3.1	Sourcing the right market data	11
3.2	The calibration process	12
3.3	Inputs disclosure	13
4	Pricing models	15
4.1	Models: what for?	15
4.2	Smile models	16
4.2.1	Local Volatility	16
4.2.2	Stochastic volatility	17
4.3	Model selection and calibration, and model risk	19
4.4	The curse of dividends	19
4.5	Correlation modelling	20
4.6	Multi-Curve interest rates models	21
4.7	Better SABR approximations	22
5	More about algorithms	23
5.1	Closed formulae	23
5.2	Binomial trees, trinomial trees	23
5.3	Finite difference	25
5.4	Monte Carlo simulation	26
5.5	Conclusion	28
6	Beyond prices	30

1 Introduction

Every crisis comes with its reminder that financial stability is fragile, forcing regulators to act in order to prevent systemic collapse. Their objective is to increase transparency and restore confidence. Financial instruments valuations' are under scrutiny as they play a key role in accurate financial statements and sound risk management.

Driven by innovation, financial markets have grown to a high level of diversity and complexity over the past decades. Making sure that a valuation reflects the best estimate of market value is not an easy task, in particular for instruments such as derivatives. The difficulties are often perceived as technical, and indeed the theory behind financial modelling requires a good mathematical and programming background. In reality, valuation uncertainty is mainly resulting from the art of sourcing the relevant inputs. Financial markets are opaque, as a lot of trading information is not public, and the large quantity of highly volatile information (trades and quotes) make this sourcing process particularly difficult. It means that providing a valuation is not just about delivering a figure but about the process and inputs to obtain it. In this respect competency and professionalism are essential.

This document is not an introductory course on financial modelling (literature about pricing models and financial mathematical theory is already abundant), nor a recipe book with valuations examples. It intends to be a general guide for users of valuations to better understand the drivers behind the process of valuation of financial derivatives. As a valuation professional organisation, APVIF's ambition is aligned with IVSC's: **"improve the transparency of the valuation processes for the users of valuation services"**.

2 Definitions

2.1 Financial instruments

2.1.1 Derivative instruments

A derivative instrument is any bilateral financial contract that specifies the future payouts (these can be cash amounts in certain currencies or physical deliveries of some basic instruments) depending on the evolution of one or more underlyings or on decisions taken by one or the other party in a specified decision space.

When there is at least one decision to be taken by a party, the derivative instrument is also called an option.

Some derivative instruments are standardized and exchanged in listed and organized markets, others are contracts called OTC (i.e. Over the Counter) securities. Due to the huge variety of derivative instruments, it is impossible to provide an exhaustive list. However we can try to list the most commonly used instruments and features:

- Forward: non-standardized contract between two parties to buy or to sell an asset at a specified future time at a price agreed upon today.
- Future: standardized contract between two parties to buy or sell a specified asset of standardized quantity and quality for a price agreed upon today with delivery and payment occurring at a specified future date.
- Vanilla option:
 - Call option: the buyer of a call option has the right, but not the obligation to buy an agreed quantity of the underlying from the seller of the option at a certain time for a certain price
 - Put option: the buyer of a put option has the right, but not the obligation to sell an agreed quantity of the underlying to the seller of the option at a certain time for a certain price
 - Digital/binary option: the buyer of a digital receives a fixed amount (i.e. cash or asset) if the option condition is satisfied and nothing otherwise
- Option exercise types:
 - American - the option can be exercised at any time during the entire life of the option

- European - the option can be exercised only at maturity
- Bermudan - the option can be exercised at specific date up to maturity
- Swap: exchange of cash flow or a stream of cash flows between two counter-parties. The amount exchanged is fixed or linked to the evolution of an observable. The typical vanilla interest rate swap involves an exchange of fixed rate against floating rate (i.e. one party pays a determined fixed rate and the other an amount linked to the evolution of a floating rate, for instance EURIBOR 6 months). The variety of swap is vast, e.g. interest rate swaps, currency swaps, commodity swaps, credit default swaps, total return swap, variance swap, etc.
- Multi-asset (in case of derivatives linked to more than one underlying):
 - Best-of - keep only the performance of the best performer
 - Worst-of - keep only the performance of the worst performer
 - Basket - weighted sum of the individual performances
 - Rainbow - weighted sum of the individual performances the weights are attributed to each performance according to its rank.
- Barrier option: classified as exotic derivative, a barrier option will either spring into existence or extinguish depending on the underlying level crossing a certain threshold, known as barrier level. A barrier options can have very different characteristics, for instance:
 - Barrier event: rules to determine how and when a barrier is defined as crossed (e.g. continuous, daily, final or Parisian type; upside or downside)
 - Barrier type : "in" (i.e. option is activated when barrier event occurs) or "out" (i.e. option is deactivated when barrier event occurs)
- Quanto/Compo options: defines how and who supports the exchange rate risk (i.e. the investor or the option seller) when the underlying(s) and the options are denominated in different currencies.
 - Quanto - the option seller takes the exchange rate risk
 - Compo - the investor takes the exchange rate risk

Derivative instruments are typically thought as being derived from simpler ones - the basic instruments.

2.1.2 Basic instruments

A basic instrument is typically standardised and exchanged in listed and organised markets. Usually highly liquid, it is reasonably easy to retrieve information regarding these instruments (i.e. market data, especially the "official" quote/price of the instrument). We identify two categories of basic instruments: ownership securities and debt securities.

Example of ownership securities:

- Equity - represents a fraction of ownership in a company
- Commodity - ownership of a defined quantity of raw materials
- Foreign Exchange - ownership of a foreign currency
- Real Estate - ownership of a real property

Example of debt securities: Bonds - instrument of indebtedness of the bond issuer (e.g. public agencies, supranational institutions or companies). They are many variations of bonds characteristics (i.e. fixed vs floating rate, coupon vs zero coupon bonds, etc.), here are some of the most common features:

- Issuer - legal entity issuing the bond
- Principal - is the amount on which the issuer pays interest, and which, most commonly, has to be repaid at the end of the term.
- Maturity - date at which the issuer has to repay the principal amount
- Coupon - cash flow paid by the issuer during the life of the bond and/or at maturity date. The amounts paid are usually a percentage of the principal and represent the interest of the bond.
- Yield - is the rate of return received from investing in the bond
- Credit Quality - refers to the probability of default of the issuer

Basic instruments have importance in this document because, aside being traded for themselves, they act as "underlyings" of derivative instruments. Thus their own value will be used as input for the valuation of other more complex instruments.

2.1.3 Underlyings

An "underlying" (or observable) is any precisely defined value that two parties can contractually agree on, when measuring its realisation at a given date/time. A typical (and most frequent) example is the official quote (in a given currency) of a basic instrument on a given organised exchange. Fixed index and rates are other examples. Underlyings don't need, however, to be always the price of financial contracts. An underlying can also be for example an (officially measured) temperature, a rainfall level, inflation, etc. Although most underlyings are numbers, they may also be yes/no choices or an element in a fixed list of possibilities. Former is for instance the default status of a company, later could be the credit quality notation of an entity by some specified rating agency.

2.2 The concept of value

Before we enter into the details of producing a valuation, we need to define what is expected from it. Usually we expect to get the amount in a given currency unit that we would receive or will have to pay for an asset or liability that we own. Let's say we hold an asset, if it is a basic instrument, traded frequently on a liquid and organised market we can look to the last traded price which would give us a pretty good indication for its value. If necessary, we can also add some clever statistical analysis to produce an even more robust value.

But there are other situations, the ones we intend to cover in this document, which are not so straightforward. For financial instruments that are not standard nor liquid, like most derivatives, financial theory comes to the rescue. Any derivative instrument can be 'modelised' as the combination of more liquid instruments or proxies providing a quasi perfect hedge. By entering a series of transaction, we can almost convert an illiquid instrument into cash currency today. So by building on financial concepts such as discounting (the value today of a future cash-flow) and arbitrage theory (probabilities and replication) we can calculate a theoretical, model based value. This approach requires some assumptions about market behaviour, liquidity... that were proven wrong. However the theory holds because with diversification and controls of risks the results are satisfactory.

This said, which output are we targeting:

1. An estimated fair value?
2. A valuation that reflects at a given time the price at which this instrument would likely trade in the market (mid-price)?

3. Or the expected transaction price between party A and party B?

For example in the first situation if we want to value a call option, we only need a Black-Scholes formula, interest rates as published in newspapers and historical volatility, in the absence of any other information available this would represent a reasonable estimate.

In situation 2, if the instrument is negotiated on a listed market with clearing this would involve changing to OIS discounting. We should also use implied forwards and volatilities calibrated on market prices for similar instruments.

In the last situation we could also factor cva-dva-fva or take into account particular collateral arrangements between parties and liquidity adjustments in our pricing.

An easy case is when similar instruments are trading on a liquid and transparent market, then clever interpolation can provide a good proxy for the valuation. So the decision to choose a particular model and calibration for a set of parameters is closely linked to the selection of instruments available in the market with enough liquidity that will be used to hedge the transaction. In some situations, a product is itself a combination of more simple components, its valuation can be obtained by valuing each of them independently.

When it is impossible to identify liquid instruments to provide the hedge, the charge for non hedgeable risks is increasing and the impact of model and parameters is less relevant. In this case, usually the choice of methodology is more a consensus approach between market participants than a rational outcome.

Most trades occur with a price maker on one side, typically an investment bank. They offer liquidity to price takers in exchange for a fee / margin and may also include a pricing buffer for non hedgeable risks. To what extent should valuations take into account these external costs including structuring fees, credit, risks that cannot be hedged, administrative (legal, regulatory) costs...?

A valuation professional must be able to adapt to any of these situations, understand the needs of its clients and clearly explain his decision to use a particular methodology.

So by definition the value carries some uncertainty; models assumptions do not represent the real world and there are residual risks that cannot be taken into account when modelling. Furthermore most model inputs include high levels of uncertainty. To overcome this problem, a solution consists in explaining and justifying the valuation with details on how it was produced and the uncertainty attached to it. A valuation is not only a number but a

detailed description of its purpose and the entire process to produce it.

2.3 Valuation methodology

The valuation methodology can be decomposed into a few steps:

1. identify product characteristics,
2. identify relevant market data for hedging,
3. selecting the most appropriate model and implementation given 1 and 2,
4. calibrate inputs for the selected model to market data,
5. check outputs for consistency.

This is a valid approach when 2 is based on very liquid instruments. Sometimes 4 requires a more continuous monitoring as data available for 2 is volatile and may include a lot of "noise". This means that 4 is part of a whole process independent from the valuation and requires a full infrastructure. It may even happen that 2 can even return nil. In this particular case, a subjective proxy based approach will be necessary. Only experience can provide a solution. In particular below is a list of some basic considerations worth mentioning when faced with a valuation assignment.

2.3.1 Do we understand the product for which we must produce a valuation?

This seems to be a trivial question but experience has proven that too often the importance of this step is underestimated. As mentioned above a financial derivative is a contract, so a legal document describing the rights and obligations of each party. It is fundamental to understand in details such rights and obligations. As with any legal documentation, a small line can make a great difference on the outcome.

2.3.2 Risk free rate?

How to assess the difference between: I will give you one tomorrow and I may give one at a future date. This seems obvious for everybody but translating this concept into numbers requires attention. Earlier literature referred to the risk free rate but more recent work implies that forwarding rates and discount rates should be chosen with care to take into account the specificities of products and counter-parties. Again this means that one solution fits all situations should be avoided.

2.3.3 Should models define inputs or inputs define models?

A common belief is that the most sophisticated model will produce a better quality valuation. This is flawed, from experience, a sophisticated model with wrong inputs may produce a far worse result than a more basic model with robust inputs. It doesn't mean that sophisticated models should be banned but the valuer must find the right balance and focus primarily on sourcing accurate inputs. A sound approach is to check if the chosen combination of model and inputs reproduces observable market prices for similar derivatives instruments.

2.3.4 Fair value or valuation outputs?

In most cases, it is important to add sensitivities or inputs impact when providing a valuation. This can help analyse discrepancies between two providers and resolve disputes. Of course disclosing data sources is also necessary, whether in the form of market reference quotes or trades or as pre-processed model inputs.

3 Market data and calibration

The purpose of this section is to:

- Describe the market information used for the production of the valuation
- Present the quality and relevance of market data, and in particular discuss the fundamental difference between quotes, trades and consensus based inputs.
- Review the different levels of observable inputs, market noise, bid ask spreads and more generally liquidity and credit adjustments.
- Describe the calibration process.
- Examine cases where there are no available market data and talk about proxy methods and resulting calibration errors.

As mentioned above, the valuation of a derivative instrument involves many steps and choices. These choices are dictated by technical knowledge, market situation and use, but also the goal of the valuation. Probably the most difficult choice is selecting the right inputs. Some of these inputs can be directly observed as market data, others (aka parameters) need to be implied from market data in a process called calibration.

In any case, selecting the appropriate market data to use is a judgement: one has to investigate liquidity of market data used and to check that it represents at best the risk factors one wants to model. This subjectivity is an inevitable side effect of the following situations:

- Financial markets information contains a lot of noise. So in a short time frame, one can observe contradictory market data. The sourcing process should therefore include filtering to extract the most relevant information and discard noise.
- At different points in time, best market data to use and thus the appropriate model to use may change, requiring adjustments to the valuation methodology over time.
- When calibration is necessary, the valuation provider has to choose a numerical method meaning a numerical precision. This will imply a trade-off between numerical quality of the valuation output and computer and time costs for achieving it.

3.1 Sourcing the right market data

Market data are market information from which the valuation provider will infer the valuation. We can distinguish different situations for market data:

- it is based on trades for which there exists a history of reliable transactions of decent sizes. In this case, the art of the sourcing is to provide the corresponding price of the transactions for the relevant time. This method is considered to be very reliable and should be preferred to the other method given the size of the trades are significant compared to the size of the position to value. A typical case is when a valuation provider is asked to value some very liquid instruments linked to large capitalisation stocks, liquid futures, common foreign exchange positions, vanilla interest rates transactions. It is important to provide bid/ask spread; the source of the transactions (timing of the data: snapshot time, market data origin). For large positions which may have a market impact, the valuation should include some additional correction to account for large transaction bias and may clearly indicate this add-on.
- A variation of this method is when available market data is based on broker quotes. In that case, as the prices used are not real transactions, but quotes, the valuation can be considered to be less accurate. This is because quotes can sometimes be quite different from prices at which a transaction would occur, some quotes are even referred to as indicative prices. Obviously, this situation should be treated with caution.
- Another variation of this method is based on consensus estimation. The valuation provider asks a consensus pool to provide their own estimation of the price for a given instrument. He/she will then use as a result the average value of the consensus with some data cleansing to remove any outlier. Additionally, the valuation provider may indicate the standard error of the contributions; explain the methodology to eliminate any outlier or questionable contribution thanks to a pricing challenge mechanism. Again, this approach is limited as the set of financial instruments on which one can get a contribution is restricted.

The worst case scenario is when there is not any reliable market data available to produce inputs, neither trades, quotes, nor consensus data. In this case, the valuation provider has to rely on a model based inputs that in turn use a set of historical data and/or indirectly linked market data. These second level model inputs can be obtained for example from similar instruments so called proxies.

As a general concern, the valuation provider must prioritise information; trades, quotes, consensus, proxies and historical data. The pertinence of market data information is critically depending on the liquidity and/or frequency, or say differently, the size and number of the recorded transactions. The larger the volume and the set of transactions, the more reliable the price information is. Listed markets are a good example of reliable information. Broker quotes, ECN, dark pool quotes depends again on the size of the market. A new methodology that is becoming more and more popular is VWAP (Volume Weighted Average Price). This incorporates somehow volume information and enables to capture market impact and should be examined in case of large market impact.

For a valuation provider, it may be critical also to do some data cleansing where information for the same financial instrument is taken from different sources and compared to identify market noise. Information about bid/ask spread may also be interesting to provide to the final user to give an estimate of how tight the market is. In addition, in valuation, a more and more common market practice is to incorporate liquidity, funding and credit valuation adjustments. This is referred to as LVA, FVA, CVA (Liquidity Valuation Adjustment, Funding Valuation Adjustment, Credit Valuation Adjustment) but also DVA (Debt Valuation Adjustment) and bilateral CVA. This new practice enables to incorporate in the valuation the cost of funding, liquidity and credit event. These new adjustments may make the valuation a more difficult process as one will require new inputs to assess the credit quality of the counter-party. This information may not always be available. One would rely on proxy method where one uses a similar credit name (based on industry, sector, region, size of the company, rating, etc...) to identify one or many comparables. As this is not standard, we recommend that the methodology and the inputs are provided to give full transparency on the computation.

3.2 The calibration process

Some model inputs known as parameters cannot be observed in the market. To obtain them, the valuation provider has to apply a process called calibration. This means reverse engineer inputs that once used with the model will price the set of relevant instruments in line with their observed market data. This calibration process can be implemented with three different approaches:

- The valuation provider maintains his own set of parameters, with a large database of observable market data that are statistically challenged with its own parameters and models in a continuous calibration

process.

- The valuation provider takes a snapshot of market data available at that point in time but does not follow regularly the markets to make sure there is no problem with market data. This approach lacks the dynamics of the first approach but should be reasonable when there are sufficient market data available and liquid underlying markets.
- The last approach is for the valuation provider to use the market parameters provided by a source. This is certainly less reliable and independent as part of the process of the valuation provider is precisely to identify relevant market data and estimate model parameters. This approach generates model uncertainty and noise especially if the model used in calibration and its implementation are different from the one used in the valuation.

This process is usually not straight forward as it relies on a finite set of market data whereas most inputs would usually have two continuous dimensions, one is the underlying level and the other one is a time horizon.

Calibration is thus the result of implementation choices:

- interpolation and extrapolation methods for parameters.
- perfect or best fit algorithm with tolerance level.
- and as mentioned before balance between numerical precision and calculation power.

3.3 Inputs disclosure

Data can be classified into different types:

- Publicly or semi publicly available information like stocks, interest rates, fx rates, futures closing or end of day prices. These data are not too hard to find and results are very similar between different sources. However, corporate actions may change the situation and make things more involved than one may think at first consideration. Also, one may observe tiny differences between different sources based on the origin of the market data and time of snapshot (true for foreign exchange rates, but even for futures, etc.). To avoid differences, it is critical to provide the source of the data.

- Professional information like volatility surface and smile, correlation, cds spreads, dividend forecast, z-spreads, ois spreads, smile information , etc... This would require specific subscription and cannot be checked easily against publicly available information over the internet. Disclosure of the source is in that case quite critical.

In a number of situations, the valuation provider needs to evaluate a product where there are no observable market data. In this particular case, one should rely on interpolation or extrapolation on comparable and disclose this information to the client. Obviously, reference to proxies is acceptable as long as the valuation provider can detail the methodology and demonstrate that the choices were done with judgement and expertise.

Last but not least, whenever a valuation provider does a model based valuation, it is crucial that he monitors the quality of the fit of the model to the market data used as inputs. Depending on the model and the data, one would try to obtain a perfect fit or a best fit. However, in all cases, the valuation provider needs to make sure he has validated somehow the calibration results. In this case as well results of this validation must be disclosed to the receiver of the valuation.

4 Pricing models

By definition, for financial derivatives, future cash flows and payouts are unknown in advance and may depend in a complex way on the evolution of underlyings and decisions taken by the parties. Mathematical finance theory, actively developed during the last decades, provides some solutions to this problem. A pricing model makes assumptions about the future evolution of one or more underlyings and decision strategies of the parties when applicable. As this evolution is typically not known today, a model makes only a stochastic assumption, assuming a whole set of possible future evolutions or outcomes. Assumptions about future evolution may depend on some parameters and/or some market data. A pricing model also defines how a price is derived from the set of possible payouts corresponding to the possible future evolutions. This is typically a kind of discounted averaging over the possible outcomes. Depending on the complexity of the model and the instrument to price, the discounted averaging may be done by an explicit simple formula (closed form solution) or involve potentially costly numerical calculations. One should distinguish between a model and its numerical implementation (see next section).

The best known and most used such proposed model is the Black-Scholes model. In this model, price calculation of simple European vanilla option on one underlying may be calculated by a closed form solution, celebrated as the Black-Scholes formula. The Black-Scholes model is however a general theoretic framework that allows for the pricing of much more complex instruments. One may want to evaluate not only a price, but also how the valuation would change should some underlying change. Valuation change with respect to infinitesimal changes are typically called sensitivities or "greeks" (the delta being the most common one) while risk analysis describe price changes with respect to bigger variations. One should keep in mind that a model is a set of assumptions on the underlying evolution depending on a set of parameters. It may be applied to many different derivative instruments.

We will not elaborate further on the basic version of the Black-Scholes model as this is now a well known and recognised model. We would prefer to focus our attention to more recent trends in option pricing models.

4.1 Models: what for?

Option pricing models cover very different situations:

- The market of the underlying is very liquid, and there is also a market of vanilla options which is also very liquid. This is the case on most equity indexes (S&P500, STOXX, SMI,..).

- The market of the underlying is very liquid, but the vanilla option market is not developed or mature enough.
- The option is written implicitly or explicitly on several underlyings, and is sensitive to joint features of these underlyings, which are hardly hedgeable.

In the first situation, vanilla options can be used effectively to hedge risk. It is therefore essential to capture the option smile in a very precise way. This is the purpose of smile models. However despite the numerous efforts devoted to build ‘the right model’, there is still no clear consensus in the market. The coexistence of many models leads naturally to model risk. Even though option pricing models are associated with volatility, this is not always the main driver. For example in equity derivatives pricing, one of the biggest issues is to forecast dividends in a sound way – this is known as the dividend curse. The same difficulty pertains to correlation modelling which is the situation described in the third point.

Two noticeable innovations in the recent years are:

- the appearance of multi-curve interest models.
- Better SABR approximations. The classical SABR model is widely used, especially to account for interest rates derivatives smiles but until recently, practitioners used to work with unsatisfactory approximations.

4.2 Smile models

4.2.1 Local Volatility

This is where Dupire’s Local volatility (LV) model (1994) has been designed to capture the smile in a perfect way. One can even show that it is completely equivalent to have a continuum of option per strikes and maturities and a LV model. There are 2 major issues with Dupire model though:

- This equivalence is of theoretical nature: in practice there is at least the issue of interpolating price across maturities, and even in some situations across strikes. There are plenty of consistent ways of performing such interpolations, and no objective way to select a preferred one. This is a big issue: typically different interpolation/extrapolation choices will provide almost the same vanilla prices for strikes close to the market ones. Yet, if one price a barrier option with a barrier which is active in a zone with a low(strike, maturity) input data density, its price will depend heavily on the extrapolation/interpolation choice.

- LV model is of static nature: it will perfectly calibrate the vanilla market at a given time stamp. ..yet at another time stamp, the ‘perfectly fitting’ local volatility will presumably change (and does so in practice).

To cope with these issues, several attempts of parametric families of LV models have been made, like Blacher’s model. The ultimate goal would have to have a very parsimonious model whose parameters have an intuitive meaning. Unfortunately no model has reached this goal so far.

4.2.2 Stochastic volatility

Stochastic volatility models (SV) are an alternate route to cope with the smile. Unlike the LV models family, they are parsimonious parametric models with parameters with a clear financial or physical interpretation. The drawback is that we lose the perfect fit property: SV models will put heavy constraints on the smile. SV models postulate a stochastic dynamic of the instantaneous volatility of the underlying. Note that this volatility is difficult to observe in practice, or to recover from daily sampled spot trajectories. The two preferred ways to estimate it is either from high-frequency data (which are available on equity indexes), or from variance swap quotes. Note that on equity indices where a very dense grid of strikes for the shortest –term options are available, it can also be inferred in a model-free way (under the assumption that the underlying has no jumps) from the option quotes.

The Heston (1993) model became the reference SV model, at least for benchmarks. The Heston model is a time-homogeneous-model, which means that its parameters do not depend on calendar time – or on the maturity. They remain therefore easy to understand. The Heston model calibrates the smile reasonably well on equity and FX. In general it behaves very well on long-dated smiles, but fails to calibrate long-dated and short-dated smiles equally well. This is not surprising since the implied volatility is very constrained: the long term instantaneous variance θ drives the long term ATM volatility but also the initial slope of the term structure of the ATM volatility... The Heston model has two other drawbacks: the instantaneous variance spends too much time close to zero, and the short-term skew remains constant when the spot and the volatility moves. This is inconsistent with the empirical observation (mostly on FX) that the skew will behave stochastically). To cope with this second effect, extensions of the Heston model have been provided. A nice one is the Double Heston model, where the instantaneous variance is the sum of two independent components with the same dynamic as in Heston.

This Double Heston model displays a stochastic skew effect, and allows also to cope separately with the long-term and the short-term smile, thus

providing a much better fit to market data. The shape of variance curves in the Double Heston model is also richer than in the Heston model, where they are monotonous. Of course there is a cost; the Double Heston model has twice as many parameters as Heston.

To cope with the first effect, the Double Log Normal model has been designed by Gatheral. A third route to fit the smile is to directly provide an arbitrage-free interpolation/extrapolation of the prices, or of the implied volatilities, in a low-dimensional parametric way. The first successful attempt in this way has been made by Gatheral and Jacquier in 2012 with the SSVI model. SSVI calibrates quite well (well within the bid-ask spread) on equity indexes. It suffers yet from some structural constraint which should be removed in the next versions of this model. Those implied volatility models have the potential to be used in a systematic manner as arbitrage-free smoothers before the usual model calibration stage. They provide also a bunch of new tools for risk management and the dynamic simulation of implied volatility surfaces. More generally, this **purely analytical** approach has been christened PEACOCKS in the academia. The challenge is to provide low-dimensional PEACOCKS, either directly in term of call prices or in implied volatilities. Each PEACOCK is a meaningful (meaning arbitrage-free) **model** for vanilla prices. It should be emphasized that we don't know yet how to associate a stochastic volatility model to a PEACOCK with adequate properties. This is a topic of active research. Exhibiting PEACOCKS with market-like shapes is another (related) research topic. SSVI is so far one of the most successful attempts in this direction.

The Local Stochastic Volatility Models (LSV) try to get the best of both worlds: a good static fit of the vanilla quotes, and nice dynamic properties. They can be seen as a local volatility perturbation of stochastic volatility models, or as a stochastic volatility perturbation of the local volatility model. The first point of view is often the right one in practice. Interestingly enough, a LSV is often calibrated through a mixed statical/historical procedure. LSV models are heavily used on FX markets, where liquid one touch option prices are available. One observes empirically that a pure LV or pure SV model calibrated on the vanillas will systematically over or under price the OTM options, whereas the LSV price is much closer to the market price. A remarkable fact is that there are few academic papers on LSV models – and there are serious issues of existence, uniqueness, and also numerical issues. This is one of the many areas where the market practice is ahead of the academia – this also means that many potential issues will pop up at some stage...

4.3 Model selection and calibration, and model risk

Many models have been proposed and several ones have industry-grade implementations. Among these industry-grade models, on a given market, many ones will calibrate almost perfectly (within the bid-ask spread). It is in general difficult to find an additional discriminating criterion in order to select the right model - the most natural one being to fit the prices of liquid exotic options, but this is a rare animal, except maybe on liquid currency pairs on FX markets.

This sets the stage for model risk. In a few words: given the consensual market information (the history of the underlying, the current vanilla prices) available, what is the range of possible prices for a given exotic option? Historically the first attempt in this direction has been made by Avellaneda and co authors with the UVM Lagrangian Model (1995). Recent advances have been made by Henry-Labordère, Touzi and co-authors, who manage to compute algorithmically or even with semi-closed formula the price range in a variety of situations. Their work is based on a remarkable application to option pricing of the known-for-ages Monge-Kantorovitch approach to optimal transport theory. This is fresh material yet; the semi-explicit formula encompasses classical variance swaps and hardly any other play-offs. The algorithmical approach is not proven, nor easy to implement. This being said, Henry-Labordère and Touzi is a breakthrough, and one should expect more and more relevant variants or refinements altogether with improved implementations in the forthcoming years.

4.4 The curse of dividends

Equity derivatives prices depend heavily on the dividends forecasts up to option maturity. On markets quoting liquid European type options (like the S&P500, even if American type options also exist), the market implicitly quotes a compounded dividend at the option maturities, which can be inferred from the put-call parity, considering that the repo is known. There are remaining issues though when an active market of dividend swaps and options is available: the assumption that the dividend is deterministic can not account in general for the available quotes. Hans Buehler has developed a very satisfactory family of models to handle jointly the vanilla quotes and the dividend swaps and options. Many underlyings (most of American single stocks and many of European single stocks) have American-type options quoted. In such situations there is no easy way to extract an implied compounded dividend from the market, and the interpretation on option prices becomes an issue: the market practice to extract an implied volatility con-

sists in using the implied volatility computed from a binomial tree but with which forwards or dividends are inputs? This first-order uncertainty should be kept in mind: an erroneous dividend assumption will produce artificial volatility smiles that the model (LV, SV or LSV) will compensate or will try to compensate so that the calibrated model parameter will produce in general erroneous exotic option prices.

4.5 Correlation modelling

Information on the dynamic of the individual underlyings of a multi asset options is generally available through option quotes; if not, historical calibrations on each of the underlying history may be achieved. In many models, the behaviour of the statistical estimators of interest for this purpose is well understood and this historical approach is often an effective one. This is not the case for the joint distribution of the underlyings – even when considering only pairwise behaviour: the notion of correlation itself turns to be tricky, and the correlation statistical estimators turn to be too noisy to yield usable figures. A noticeable improvement in correlation modelling has been the use of Wishart processes to model stochastic joint distributions. Yet many estimation and correlation issues do remain.

To cope with these issues, 2 families of models are used in practice:

- local correlation models, which allow to get a refined view of the correlation with respect to the constant Black-Scholes model. Those models are used, for instance, to solve the FX cross smile puzzle: specifying a joint distribution of 2 pairs of currency pairs with the same base currency such that the cross pair smile, which is observed on the market, can be reproduced. Guyon has proposed a very nice non linear model to get iteratively local correlation coefficients which fit the cross-smile. Note that there is still model risk remaining since several fitting local correlations can be constructed.
- Avellaneda UVM model, with an uncertainty on the correlation matrix. This approach is very relevant in many situations, due to the scarcity of information on the joint distribution. Yet it is a challenge from a numerical point of view, due to a high degree of non-linearity in several dimensions. Guyon and Henry-Labordère proposed an innovative Monte-Carlo type algorithm which is easy to implement and gives satisfactory results, at least in small/medium dimensions. Another innovative approach is the Randomized Control approach created by Huyen Pham, which yields efficient numerical implementations for the same kind of problems.

4.6 Multi-Curve interest rates models

In the post-crisis interest rates world, the interbank curve is not unique any more. In fact after the crisis, a large basis spread has appeared between the swap rates with the same maturity, but based on Libor rate with different tenors. Likewise, deposit rates and OIS rates for the same maturity present a spread of few basis points. All these spreads were often assumed to be zero when constructing zero-coupon curves. Also for reasons of liquidity and credit risk, the conventional relationship between rates of different tenors (ex. Composition of 3-month rate to deduce the 6-month rate) are no longer valid. Market participants are faced with several questions: First methodologies issues:

- What “risk-free yield curve” used to discount cash flows.
- How to build forward curves ?

But also issues of modelling:

- How to price Caps, Floors and Swaptions options consistently ?
- How to price exotic products and measure their risks ?

Various approaches were proposed to model interest rate market post 2008:

- The multi-curve approach Bianchetti (2008): Bianchetti introduced for each tenor an adjustment factor to correct the relationship of classical no-arbitrage between forward rates and zero coupon. This factor measures the difference between the forward curve and the basis curve (discounting). The pricing of standard instruments requires the calculation of this factor. Bianchetti considers that the forward curve corresponds to a foreign rate and the discounting curve to a domestic rate. He proposes a log-normal dynamic to forward Libor rate and exchange rate and proves that adjustment factor depends on their volatilities and correlation between them.
- The HJM-LIBOR Moreni and Pallavicini approach (2010): In reality there is not enough option prices, most of them correspond to limited tenors. So each extension of classic model must take into account the lack of market information in calibration. In the paper, Moreni and Pallavicini propose an extension of the HJM model which incorporates all yield curves. As in the LMM, the proposed model specifies directly the rates diffusion observed on the market, while providing beneficial Markov properties for the simulation and calibration.

- The collateralisation approach Fuji, Takahashi, Shimada (2009): These authors present a framework for modelling the collateral with an explicit modelling of the rate of account of collateral, and propose an HJM dynamic to price options. This work is very interesting, but it is enough heavy and theoretical to implement in practice.

4.7 Better SABR approximations

The SABR model is popular among practitioners specially for modelling interest rate volatility smile. It provides more flexible and an intuitive parametrisation of volatility smiles. The Hagan asymptotic expansion of the Log-normal implied volatility for SABR model, presents some problems. This formula becomes inaccurate and can imply negative density near to the forward, when the volatility-of-volatility is large or the CEV exponent is small. New accurate expansions are developed recently. Antonov & Hugu (2013) proposed, when the correlation is zero, an exact formula for the price of an option. They proposed also an approximation for the general correlation case. Balland (2013) present new representation of volatility as a normal SABR implied volatility with absorption at zero, and develop an asymptotic expansion for it and pricing formula for normal SABR. The results of such representation of implied volatility are shown to be more accurate than the original SABR expansion.

5 More about algorithms

A common misconception is to confuse pricing models and their implementation. How often have we heard, we use a Monte Carlo model. What we are going to describe in this section is which implementations are available for model resolution and how some computer algorithms development details can affect the outputs. In many implementations, there is often a choice at some point that will affect accuracy:

- number of steps in binomial tree
- random number generator and number of simulations in Monte Carlo
- number of points or resolution method in a finite difference

Our objective in this section is not to describe in details all these issues but to make sure our readers understand they exist and can have an impact if not managed adequately on valuations accuracy.

5.1 Closed formulae

Closed formulae are the perfect mathematical expression of the solution of a pricing model. As such, it is easy to implement and very fast. The most popular closed formulae is the well know Black Scholes formulae used to price European calls and puts. However generalisation to more complex pay-offs or models is difficult. In practice, closed formulae are often used to teach pricing models or to benchmark more complex implementations.

5.2 Binomial trees, trinomial trees

Sometimes called the binomial model (since it may be considered either as a model for the behaviour of the underlying asset or as a numerical procedure for the solution of the differential equation that option prices satisfy), the binomial tree method was first introduced by Cox, Ross and Rubinstein in 1979, and has been very popular among practitioners in options pricing for years (although it has been surpassed since by more sophisticated modern numerical methods). The reason for this popularity lies in the fact that this method is simple, intuitive, and relies on basic absence of arbitrage opportunities considerations rather than complicated mathematics.

Performing a binomial tree valuation involves dividing the life of the option into a large number N of time steps of length δt . We then assume, for each time step that the price S of the underlying asset moves up or down by

a specific factor u or d with the probabilities q and $1-q$. The values of the asset price should be $S_u = uS$ and $S_d = dS$. The parameters u , d and q are chosen to match the drift and variance of asset price changes over a period of one time step. In their paper, CRR assigned the following values for a log-normal process approximation: $u = e^{\sigma\sqrt{t}} = \frac{1}{d}$, but other alternatives are possible under the same constraints. This process, when repeated N times leads to the construction of a recombining tree of asset prices from the initial asset price S_0 at time 0 to $N+1$ asset price values at the expiration date. Let $O = O(S, t)$ be the value of the option at time t and asset price S and $O_u = O(S_u, t + \delta t)$, $O_d = O(S_d, t + \delta t)$ be the values of the option at the end of the period. The authors showed that we can build a risk-less portfolio P containing the option and a number $-\Delta$ of the underlying asset. if r is the constant risk free interest rate over the time period, the value of the portfolio at time $t + \delta t$ should be $P(t + \delta t) = P(t)e^{r\delta t}$. Solving the equations $O_u - \Delta S_u = O_d - \Delta S_d = e^{r\delta t}(O - \Delta S)$, we have:

$$\Delta = \frac{O_u - O_d}{(u - d) \cdot S}$$

$$O = e^{-r\delta t} (pO_u + (1 - p) O_d), \text{ with } p = \frac{e^{r\delta t} - d}{u - d}$$

These results illustrate the ideas of hedging and risk-neutral valuation as they show in an intuitive way that the price of an option does not depend on the probability q of the underlying asset and can be interpreted as its discounted future value in a risk-neutral world. Starting at the end of the tree (the expiration date), we know the values of the option that are given by the pay-off function. Going backward, we calculate the values of the option at each node for the preceding time steps until reaching the root of the tree (current asset price at the current date). We can even control these values for a potential early exercise for American options or other features (convertible bonds conditions, barrier crossing, ...). Binomial tree have been very useful for valuing options with no available analytical solution, and they still are a good introduction to understand the basics of derivatives pricing. However, as numerical procedures, they are not very flexible and their convergence is in some cases problematic. Over the years, they have been improved, enriched, fine-tuned and sometimes twisted in order to adapt to more complex models or more exotic options. A better alternative is to switch to trinomial trees that offer more flexibility than binomial trees. Here, each node has three successors instead of two. The arbitrage considerations are not valid any-more, but the price is still the discounted weighted average of the three succeeding prices. The tree can be calibrated in order to match

the first moments of the underlying's continuous process and the additional degrees of freedom may be used to adapt the geometry of the tree to more complex situations regarding the underlying process or the option contract features.

5.3 Finite difference

The finite difference methods solve the differential equation that drives the option price using an approximation of its derivatives on a grid (X_i, t_j) where X is the asset value (or a function of the asset value) and t is the time. We can write the difference equation at times t_{j+1} or t_j . In the first case, the method is called the explicit finite-difference method, since the option values at time t_j depend only on the option values at t_{j+1} . In the second case, the method is called the fully implicit finite-difference method.

We can also take the average (weighted with a parameter θ) of the two evaluations. This is called the θ scheme finite difference method.

A special case for this method is when $\theta = 0.5$ which is best known as the Crank-Nicholson method. When $\theta = 1$ or $\theta = 0$, we are back to the explicit or the implicit methods.

Before solving this problem, one has to set the initial conditions (or the final conditions since we solve this problem backward) and the boundary conditions. Depending on the features of the option, the boundary conditions can be represented by three different types:

1. Type 1: A given (or pre-calculated) value at the boundary. This type can be used for a constant barrier option, assuming the barrier is on the extreme point of the grid, or for most vanilla contracts for which the value can be precomputed for small and large values of the asset price.
2. Type 2: A given slope of the option value at the boundary. When pricing a time dependent barrier option, it is not always possible to make the barrier match up with the grid. An accurate approximation would use a linear interpolation between the two points surrounding the barrier. For other options, we may know the slope of the option value for small and large values of the asset price rather than the option value.
3. Type 3: When the option pay-off is linear in the asset for small and large asset values, it may be easier to set the boundary condition as: $\frac{\partial^2 V}{\partial X^2} = 0$

Once this step is performed, the problem can be written (except for the explicit method which is simpler) as a Matrix equation $M_j V_j = R_{j+1}$ where

M_j is a tridiagonal matrix, and V_j and R_{j+1} are column vectors representing the option value and the right-hand side of the equation. It is inefficient and very (computational) time consuming to invert the matrix M_j itself. Instead, one can perform an LU decomposition which consists in writing the tridiagonal matrix as a product of a lower triangular matrix L and an upper triangular matrix U. Computing the solutions for each matrix is very fast and accurate. Another way to solve the problem is to apply an iterative method such as the Jacobi method or the SOR method (successive over-relaxation). These methods are slightly slower than a direct method, but they are more adapted to American options.

The finite-difference methods may seem harder to manipulate than tree methods. Yet all the techniques described here are well described in academic papers and books, and since these methods are applied in other areas than finance, one can find more material to improve the convergence or adapt the method for new models or complex products. Indeed, these methods are very flexible in handling different models (from the general form of the Black-Scholes model to stochastic volatility, jump models or stochastic local volatility models), different forms of dividend payments and a large set of exotic products. The explicit finite-difference method is similar to the trinomial tree method. It is easy to program for both European and American options, flexible and copes with time and asset dependent parameters. Its main drawback is that it does not converge unconditionally. The Crank-Nicholson method on the other hand requires more care in its implementation but is unconditionally convergent and may be much faster than the explicit method for a given accuracy despite its relative complexity.

5.4 Monte Carlo simulation

The use of Monte Carlo simulation in finance is a result of the representation of a derivative price as the expected present value of its payoff. Valuing an option price by Monte Carlo involves generating paths of the stochastic process governing the asset price and evaluating the pay-off function along these paths. An estimator of the option price is given by the discounted average of these values. The estimation error is equal to $\frac{\sigma_P}{\sqrt{N}}$, where σ_P is the standard deviation of the pay-off function from N simulations. It is that simple! At least in theory.

In practice, generating paths of a stochastic process requires a lot of care.

First, the random numbers generated by computers are not genuinely random but rather produced by completely deterministic algorithms. A random number generator, often called *pseudo-random* number generator, produces a finite sequence of numbers. This means that the sequence will repeat it-

self if we simulate a very large number of pseudo-random variables. A good generator should therefore have a large period and produce numbers that are hard to distinguish from real randomness. A typical pseudo-random number generator is known as Mersenne Twister. It generates uniformly distributed numbers. In order to obtain normally distributed numbers, a common implementation is to use the Box Muller transform.

Secondly, generating paths of a stochastic process may be less straightforward as it may seem. Some processes require more mathematics and more computation than the special case of the geometric Brownian motion.

Finally, some continuous-time stochastic processes may not be simulated exactly at a discrete set of dates. Discretisation errors may in some cases change the probability distribution and produce large errors in the simulation. Taking a small time step reduces the error but it is time consuming. It is possible in some cases to reduce the error by changing the "numéraire", but with complicated stochastic processes it may be advisable to use more refined discretisation schemes. Fortunately, all these topics are well documented in the literature and may be tackled on a step by step basis.

The convergence rate of Monte Carlo methods is relatively slow. However, since this rate holds even for problems in d dimensions, these methods are competitive with alternative methods in many cases. The other advantages of these methods are generality (complex models can be handled as long as we know how to simulate the paths) and flexibility (complex pay-offs, path dependency may be taken into account quite easily). The main drawbacks are twofold.

- it is computer intensive due to the square-root rate of convergence. This can be partially solved by variance reduction techniques that are aimed to reduce the constant σ_P in order to speed up the convergence, or by using quasi Monte Carlo methods that provide a faster convergence rate.
- Monte Carlo methods in their basic form are limited to value European options. Pricing American options requires knowing the option prices for all asset values at any time, whereas simulation allows only an estimation of the option price at one point. Several algorithms have been developed to address this problem. They all require a substantial computational effort and are much more sophisticated than the basic simulation. The most popular among these methods is that proposed by Longstaff & Schwartz and it combines the simulation of forward paths at discrete times and an approximation of the value of holding the option at time t by a linear combination of known functions defined by the user.

Variance reduction techniques allow in some cases substantial gains in efficiency mostly by taking advantage of specific features of a problem and are well described in the literature. We will just briefly describe two broadly applicable methods:

- Antithetic Variates introduces negative dependence between pairs of random numbers in order to reduce the variance without changing the law of the simulated process. For a symmetric distribution, the antithetic pair will contain a first input and its symmetric. For example for a uniformly distributed variable over $[0,1]$, the pair would be u and $1-u$, and for the Normal distribution we would choose z and $-z$.
- Control Variates exploits information about the estimation error of a known variable to reduce the estimation error of the option price. To illustrate the idea, suppose that we want to price an Asian call option A . We can choose as known options vanilla calls C_i , $1 \leq i \leq p$ with the same strike price and with maturities equal to the averaging dates. We first run n simulations to value the Asian option and the known options. With this data $(A_j, C_{i,j} \text{ for } j = 1, \dots, n)$, we perform a regression between the Asian option values and the vanilla option values in order to obtain an estimate of the optimal weight ω_i for each C_i . Finally we run more simulations to estimate the price of $A - \sum \omega_i (C_i - \mathbf{E}[C])$ which is also an estimator of the price of A , with a variance that is substantially lower than the variance of A .

5.5 Conclusion

We have presented here an overview of the most commonly used numerical methods in options pricing. This list is far from complete and there is an abundant amount of literature dealing with these techniques. But the field is also in constant progress for at least three reasons.

- The computational power is still exponentially improving due to faster processors and a better use of the computing resources (grid computing, GPU programming). Techniques that were considered too slow or even impracticable a few years ago are gaining more traction.
- More and more brilliant people from other branches (theoretical physics or mathematics, applied mathematics) are attracted to quantitative finance. They bring a theoretical background and numerical tools that were almost unheard of in computational finance.

- Complexity in finance is still alive! The complexity of financial products has certainly stalled but the post crisis environment (new regulatory requirements, unusual market conditions, credit, liquidity and funding issues) creates new challenges that in turn require new solutions.

6 Beyond prices

Given all the choices and the diversity of market data sources offered to valuation professionals, it is clear that the outcome of the valuation process of financial instruments may generate some discrepancies. It would even be suspicious that two different providers come up with the same valuation for a very complex or illiquid derivative instrument. But it doesn't mean that we should be satisfied with this situation.

First of all, professional valuers should not be scared to be challenged, on the contrary. Providing a different valuation from a market participant or from another professional valuer is not necessarily a sign of incompetence. Even a liquidity provider can be wrong, mis-prices happen, the past tells us that it may take a while before a remark is identified because of using a flawed model. Even the most active players have a process of price verification as part of their market risk analysis. As discussed earlier, a lot of valuations parameters influencing the valuation require judgement as full information is not available to all participants.

This uncertainty requires a professional valuer to produce more than just a number. Transparency about the methodology, model and data used in the valuation process helps to identify the causes of the discrepancies. The user of valuations should also use different sources and assess pricing uncertainty. Beyond prices and theory, the investigation process to reconcile valuation differences is the new focus for improvements. One cannot guarantee accuracy or robustness in the output, however understanding and estimating uncertainty in a valuation is becoming compulsory. This will help to close the gap between different sources or between various types of valuations: fair value for accounting, prudent valuation and exit price.